(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) Publication number: 0 510 834 A2

(12) EUROPEAN PATENT APPLICATION

(21) Application number: 92303126.4

(51) Int. Cl.^{5}: G06F 15/419

(22) Date of filing: 08.04.92

(30) Priority: 22.04.91 US 688228

(43) Date of publication of application: 28.10.92 Bulletin 92/44

(84) Designated Contracting States: DE FR GB

(71) Applicant: International Business Machines Corporation
Old Orchard Road
Armonk, N.Y. 10504(US)

(72) Inventor: Burke, William Theodore
920 North Citadel
Tucson, Arizona 85748(US)
Inventor: Day, Kenneth Fairclough
730 North Lazy J. Way
Tucson, Arizona 85748(US)
Inventor: Dewey, Douglas William
8602 East Old Spanish Trail No. 14
Tucson, Arizona 85710(US)

(74) Representative: Blakemore, Frederick Norman
IBM United Kingdom Limited Intellectual
Property Department Hursley Park
Winchester Hampshire SO21 2JN(GB)

(54) Method and apparatus for accessing data recorded as linked lists.

(57) A WORM disk stores data in continuation chains. Each chain being arranged as a plurality of groups of contiguous data-storing areas (such as disk sectors, clusters of sectors, tracks and the like). The groups are separated on the disk. Each recorded area in the chain has a forward pointer to a next succeeding area of the chain, whether such succeeding area is in the same chain or in a next chain. A last succeeding one of the groups includes unrecorded allocated area(s) of the chain. To find the end of the chain, only the last area of each successive group is accessed and read until a last area of a last group is accessed with an indication that such last group area is unrecorded. Then the entire last group is scanned to find the last recorded area of the chain.

The present invention relates to data storing systems wherein the data is stored in linked lists, such as in continuation chains of write-once read-many (WORM) media.

WORM media have been used for years to store information in an optically-sensible form, such as by ablative recording, phase change recording, dye recording, and the like. Most WORM media employ a single spiral track; one can treat each circumvolution of the spiral track as an addressable track; such treatment emulates the spiral track to concentric parallel tracks of a disk. The present invention is practised with either concentric parallel tracks, with the spiral track or rectilinear tracks. It has been the practice to arrange the data on such media in linked lists of disk media addressable sectors commonly referred to as continuation chains. A directory means enables accessing any one of a multiplicity of such continuation chains. In fact, directory structures stored on such media are also stored in continuation chains. Such continuation chains are transparent to a computer user having a computer system employing such media.

To access the most recently recorded data in a continuation chain, it is first necessary to find the addressable area of that chain which was last recorded. Such last addressable area (sector or a cluster of sectors in an optical disk, for example) contains the most-recent data. Except for the end of a chain, no continuation chain has any null or empty addressable areas within the chain. The end of the chain is therefore indicated by a first null or empty addressable area in the chain, as further explained below.

Each of the addressable areas includes a forward pointer indicating the address of the next addressable area in the chain. In creating or extending a continuation chain, a group of addressable areas of the WORM medium are allocated to the continuation chain. Such groups of addressable areas can be of any size and can vary in size within a chain. To most easily practice the present invention, it is preferred that all groups within a chain be of a constant size, no limitation thereto intended. Each constantsized group could occupy one track of a disk medium; of course, a group may require two or more tracks, one-half track or an arbitrary number of addressable areas or sectors. When recording data in one of the addressable areas in a continuation chain within an allocated but un-filled group of such areas, no additional allocation of addressable areas is needed; only when a current group of addressable areas is filled will an additional group of areas be allocated to a chain. From this statement, it is seen that all addressable areas in a group are contiguous on the medium while the groups may be spatially displaced.

In accessing data stored as a continuation chain, it is often desired to first access the last-recorded data, i.e. find the end of the continuation chain. The practice has been to read the entire group of addressable areas into a random access memory, then analyse the contents of the read addressable areas for finding the last-recorded area (the recorded area immediately preceding a null area is such last-recorded area). A first read failure signals a null or unrecorded but allocated area; i.e. the end of the chain. Accessing the entirety of the chain is time-consuming and therefore, a faster and more efficient linked-list end-finding method and apparatus is desired.

According to the present invention there is provided a method of accessing data from a data storage means in which the data has been recorded as a linked list, the data occupying chained storage areas in which each storage area has forward pointer to the next in the chain, the chain being subdivided into groups of storage areas, the method comprising, in a first phase, searching through the last storage area only of each group along the chain to identify the first group of which the last storage area contains no data, and in a second phase, serially reading through the storage areas of the group identified in the first phase to find where the last recorded storage area is in the chain.

Further according to the present invention there is provided data storage apparatus for accessing data from a data storage means in which the data has been recorded as a linked list, the data storage means having chained data storage area in which teach storage area has a forward pointer to the next in the chain, the chain being subdivided into predefined groups of storage areas, the apparatus comprising means operable, in a first phase of an access operation, to search through the last storage area only of each group along the chain of storage areas on the data storage means to identify the first group of which the last data storage area contains no data, and further means operable, in a second phase of an access operation, to read serially through the storage areas of the group identified in the first phase to find where the last recorded storage area is in the chain.

In the drawings:-

Fig. 1 is a block diagram showing a computer system in which the present invention is advantageously employed.

Fig. 2 is a diagrammatic showing of the format in an addressable area used in a continuation chain used in the Fig. 1 illustrated system.

Fig. 3 is a schematic showing of a plurality of continuation chains and their respective access and scanning as used in the Fig. 1 illustrated system.

Fig. 4 diagrammatically shows a disk-shaped record medium usable in practising the present invention in the Fig. 1 illustrated system.

Fig. 5 is a simplified machine-operations chart showing practising the present invention in the Fig. 1 illustrated system.

Fig. 6 is a diagrammatic illustration of a second format of an addressable area for practising the present invention.

Referring now more particularly to the appended drawing, like numerals indicate like parts and structural features in the various figures. Host processor(s) 10 are connected to one or more microprocessor controlled peripheral controllers 11, these two units plus the electronic circuits in the disk devices 12, 13 constitute computer means for operating the devices. Controller 11 in turn is connected to one or more optical disk devices, herein termed players, 12, 13. Players 12, 13 preferably also record data on disk media. In many optical players, the media are removable. In a best mode of the invention, the optical media are WORM media. Data recorded on disk media usable in players 12, 13 are recorded in continuation chains, as later described.

Each WORM disk medium typically has a multiplicity of tracks, each of the tracks being divided into a plurality of addressable sectors. In the illustrated embodiment, each sector constitutes a respective one of the addressable areas. The format of each addressable area 15 includes a reverse pointer 16 pointing to the addressable area which was immediately previously recorded on the disk medium; such a reverse pointer is not necessary for practising the present invention. A data field 17 stores data. Such data may be user data, directory data, disk status data, and the like. In fact all data stored on a disk may be stored in this format. Finally, forward pointer 18 points to a next one of the addressable areas in the continuation chain; the last one of the addressable areas having field 17 filled with data points to a next addressable area which is null, i.e. stores no pointers 16,18 nor does it have a data field 17, or has a datafield which is either empty or partially filled.

Fig. 3 illustrates two continuation chains 20 and 21, each chain consisting of predetermined-sized allocated groups of addressable data-storing areas 15. Accessing either of the chains 20 or 21 for finding a last entry in one of the continuation chains is next described. Directory 19, implemented as a hash table, (on disk medium 52), which itself is either constructed using continuation chain(s) or is pointed to by entries in a continuation chain, stores addresses of respective first entry points of a plurality of continuation chains; directory 19 may also point to intermediate ones of the groups for entering the continuation chain at any one of a plurality of the groups. Such chains may store sub-directory data, user data, or other control data in the respective addressable areas collectively enumerated 15.

Continuation chain 20 has allocated constant-sized groups 25-28 while continuation chain 21 has allocated constant-sized groups 40,43,45 and 49. Groups 28 and 49 respectively contain the last entry of the two continuation chains, respectively. The prior art method of finding the last entry of these groups 28 and 49 is to read the entirety of all preceding groups of the respective chains into a memory for analysis. Such extensive reading results in unnecessary elapse scan time for finding the last-recorded area; reading only the last addressable area of each group in a chain more quickly and efficiently finds the last entry of any continuation chain.

The addressable areas 15 in first group 25 are recorded serially beginning with addressable area 21A, this addressable area stores the first entry of the continuation chain. The intervening addressable areas 15 are then recorded in seriatim to the end of group 20, thence in later-allocated groups 26-28 until the last recorded addressable area 30 has been recorded. The remaining addressable areas 34, 34A of group 28 are unrecorded allocated addressable data-storing areas. The searching for the last recorded area 30 of continuation chain 20 consists of reading only the last areas 31, 32, 33 and 34A of each respective allocated group of the chain. The searching is completed by first identifying the unrecorded allocated area (null) 34A; all prior examined last areas 31, 32 and 33 in the preceding groups of the chain were recorded and therefore do not signify the end of the chain. Null area 34A indicates that the last recorded area is either in allocated group 27 or 28. The next step in the search is to attempt to read the entirety of allocated group 28 into host processor 10 for analysis; the recorded into and allocated area 30 of group 28 having the highest address is the last or most-recently recorded data of the continuation chain. A read failure in the first unrecorded area 34 identifies area 30 as the last recorded area of chain 20.

The last recorded area of chain 20 at one time was area 33 of group 27. When area 33 of group 27 was to be recorded, then group 28 is allocated but not recorded into. Then area 33 is recorded with a forward pointer 18 to allocated but entirely unrecorded group 28. Accordingly, if the allocated areas of group 28 are loaded into host processor 10 and if all of the areas are still unrecorded, then the last or most-recently recorded area is indicated to be area 33 of group 27. To facilitate this type of searching, it is desired to have the allocated groups have the same number of addressable data

storing areas.

A second continuation chain including allocated groups 40, 43, 45 and 49 show interleaving of allocated groups between two continuation chains. In a practical embodiment, a large plurality of unrelated continuation chains have allocated groups interleaved; many of such continuation chains may have differing number of data-storing areas in each group. As an example of a latter mentioned continuation chain, certain user data to be immediately recorded will be allocated a number of data storing areas 15 equal to that needed to record data leaving no unrecorded allocated data storing areas. All of these chains are interleaved. This interleaving means that the circumferential beginning of allocated groups vary. However, with constant-sized allocated groups, the beginning address of a next constant-sized allocated group is indicated by forward pointer 18 of the last area of each allocated group. Since the size of each group is apriori information, such information can be a part of the programming of host processor 10 or separately stored as in a register or storage location of random access memory 14, the address of each last area of a next constant-sized group is calculated by merely adding the number of areas in such constant sized groups thereby enabling a fast search of the contents of each last allocated area in each allocated group in the continuation chain. In this regard, the apriori information is known to host processor 10 and is an indication of the displacement of each groups last allocated area from the first allocated area of each such group. Accessing each last area of any group from the previous last area of a preceding group is by adding the number of areas in the group to the value of the forward pointer 18. In starting the scan of last areas in each group, the address of area 21A of continuation chain 20 found in directory 19 is modified by the number of areas in group 25 for immediately accessing last area 31 of group 25.

Fig. 4 is a simplified diagrammatic plan view of a WORM disk 52 insertable into and removable from any of the optical players 12, 13. A copy of the directory 19 is stored beginning at one radial extremity of disk 52, as the directory size increases, additional sectors of disk 52 are allocated for directory usage. Directory 19 or any chain storing control information may store the number of addressable areas in the continuation chains having a constant number of areas. Three tracks 55, 56 and 57 are linked as a continuation chain having three track-sized groups in the continuation chain, such as groups 25, 26 and 27 of continuation chain 20. The fourth group 28 is not shown. As usual, disk 52 is hard sectored into by circumferentially spaced-apart sector marks 53; the recordable areas between any two of the adjacent sector marks 53 are termed a sector which is addressable in a usual manner, i.e. sector marks 53 contain address information of sectors immediately adjacent the marks 53, respectively. Each sector is one of the addressable data-storing areas in the Fig. 3 illustrated continuation chains. Directory 19 is preferably recorded on disk 52 in a usual manner.

Fig. 5 is an abbreviated machine operations chart for the Fig. 1 illustrated system. In accessing continuation chain 20, in a first phase of operation, directory 19 is first accessed to identify the location of the first address able area 21A in chain 20. In step 65, disk player 12 scans its lens system to seek the last sector of group 25. Since the address of area 21A is known, host processor 10 calculates, as a part of step 65, the address of last area (sector) 31 of group 25. This calculated address is the number of areas in group 15 plus the address of area 21A found in directory 19. Once the seek is complete, only last area 31 is read at machine step 66 for analysis by host processor 10. At branch step 67, host processor 10 determines whether or not the last area 31 or is a null (unrecorded) allocated area. In this instance, area 31 is recorded. Host processor 10 then in step 70 reads forward pointer 18 of area 25. Steps 65-70 are repeated for last areas 32 (group 26) and 33 (group 27). Upon trying to read last area 34A of group 28, recorder 12 cannot read the area (no data). This failure to read data is transmitted to host processor 10 and indicates the end of the first phase of operation. Such failure is interpreted by host processor 10 that the area 34A is a null area; from step 67 after accessing area 34A in steps 65 and 66, machine operations proceed to step 69 in a second phase of operation wherein host processor 10 knows the last recorded area of continuation chain 20 is in either group 27 or 28. Since the probability is that the last area 30 is in group 28, all areas of group 28 are attempted to be read. Upon attempting to read the area 34 immediately adjacent recorded group 30, a read failure is signalled to host processor 10 by recorder 12. The read failure indicates to host processor 10 that the last area of continuation chain 20 is area 30. From step 69 other machine operations are performed, such as a next data accessing operation.

As mentioned above, the true last area of chain 20 could be the last area 33 of the immediately preceding group 27. In that instance, a read failure of all areas of group 28 indicates the last recorded area of chain 20 would be last area 33 of group 27.

The time-saving advantage of practising the present invention arises out of eliminating reading all areas of each group for analysis; such additional reading requires substantially greater elapsed time; in some tests such additional time when using the prior method of analysis rather than the present

invention was a factor of four-to-one and up to ten-to-one increased scan time.

The invention can be easily practised with variably sized groups. Referring to Fig. 6, addressable area 15 is modified to be addressable area 15A. Pointer fields 16 and 18 remain unchanged. Data field 17A is slightly smaller than data field 17 to accommodate new field NEXT SIZE 80. Field 80 stores the number of sectors allocated in the group pointed by forward pointer 18. In practising the present invention as explained with respect to Fig. 5, instead of adding the constant number of areas to the address indicated by the forward pointer, the contents of field NEXT SIZE is added to the address of the first area of a next allocation group. In the areas of each group other than the last area, field NEXT SIZE 80 may contain an indicator (invalid address) showing that the next area of the chain is in a contiguous next area or sector.

While the invention has been particularly shown and described with reference to preferred embodiments thereof, it will be understood by those skilled in the art that various changes in form and details may be made. For example, it should be noted that peripheral controller 11 may be a pluggable circuit board in a host computer, be embedded as a part of a host computer, an attachment card to a host computer, or be a separate standing unit connected to a host computer. Also, such peripheral controller may be implemented with the aid of computer programming in a host processor.

Claims

1. A method of accessing data from a data storage means in which the data has been recorded as a linked list, the data occupying chained storage areas in which each storage area has forward pointer to the next in the chain, the chain being subdivided into groups of storage areas, the method comprising, in a first phase, searching through the last storage area only of each group along the chain to identify the first group of which the last storage area contains no data, and in a second phase, serially reading through the storage areas of the group identified in the first phase to find where the last recorded storage area is in the chain.

2. The method set forth in claim 1 wherein the number of storage areas in each group is a constant predetermined number and

   each successive last storage area of the respective groups is addressed from an immediately preceding group by adding the constant predetermined number to the contents of the forward pointer in the last area of the immediately preceding group and using the sum of such addition as the address of the next area to be accessed.

3. The method set forth in claim 1in which

   each last storage area of each of the groups in a continuation chain, has a recording of an address pointer to a first area of a next succeeding group of the chain and an indication of the number of areas is said next succeeding group.

4. The method set forth in claim 1 comprising

   generating and storing an indication of the number of storage areas for all of the groups in the chain;

   storing an address pointer to a next area in each of the storage areas in the chain; and

   adding the said number to the address pointer of each last area of the respective groups and addressing a last area of the immediately succeeding group using the added sum as the address.

5. The method set forth in any one of the preceding claims wherein

   the data storage means is a record member and the linked list of data constitutes one of a plurality of such linked lists recorded on the records member, the method including a preparatory step of selecting one of the plurality of linked lists to access the data therefrom.

6. The method set forth in claim 5, wherein

   the recorded record member is a write-once read many data storing disk having a multiplicity of concentric tracks, each track having a predetermined number of addressable sectors, and each of said areas is constituted one of said sectors; and

   said groups have either an integral multiple or sub-multiple of the number of the sectors in said concentric tracks.

7. Data storage apparatus for accessing data from a data storage means in which the data has been recorded as a linked list, the data storage means having chained data storage area in which teach storage area has a forward pointer to the next in the chain, the chain being subdivided into predefined groups of storage areas, the apparatus comprising means operable, in a first phase of an access operation, to search through the last storage area only of each group along the chain of storage areas on the data storage means to identify the first group of which the last data storage area con-

tains no data, and further means operable, in a second phase of an access operation, to read serially through the storage areas of the group identified in the first phase to find where the last recorded storage area is in the chain.

8. A data storing system having a record member storing data in a plurality of continuation chains and a directory means on the record member indicating a first entry point of the respective continuation chains, each of the chains including a number of addressable data storing areas of the record member which are addressably linked by address pointers, the areas being arranged in serially linked groups of the areas, the areas in a group being contiguous on the media, each of the groups having a first and last one of the areas which are serially linked beginning at the first one and ending at the last one of the areas and the last one of the areas having an address pointer to a first one of the areas in another one of the groups in the continuation chain, the system including, in combination:

computer means operatively coupled to the record member for reading data stored on the record member, the computer means having data accessing program indicia for enabling the computer means to read a last one of the addressable areas in successive ones of said groups until finding a last one of such addressable areas that contains no data in a predetermined one of the groups; and

indicate program indicia in the data accessing program indicia for enabling the computer means to indicate that the addressable data having the highest number of intervening addressable areas to the single starting point is in the predetermined one of the groups.

9. The system of claim 8, further including, in combination:

said data accessing program indicia further having scan program indicia for enabling the computer means to operate after the indicate program indicia were executed to serially read the areas in the predetermined one of the groups beginning with a predetermined one of the areas and proceeding through the linking toward the last one of the areas in the predetermined one of the groups, and to detect that a given one of the areas in the predetermined one of the groups contains no data, then to stop further reading and indicating that the area in the predetermined one of the groups immediately preceding said given one of the areas in the linking contains the data in the area having the most number of intervening areas to the predetermined one of the areas.

10. The system of in claim 9, further including, in combination:

random access storage means in the computer means for receiving and storing data from the areas being read;

the data accessing program indicia including analysis indicia for enabling the computer means to analyse the data stored in the random access storage means for identifying the last data in a continuation chain.

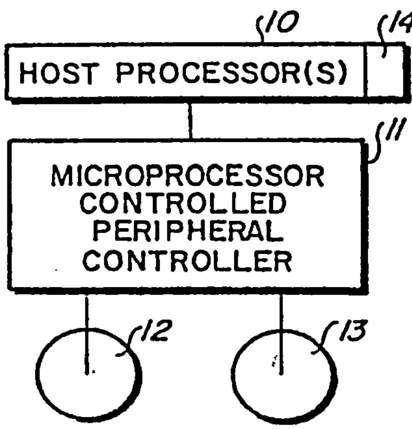

FIG. 1

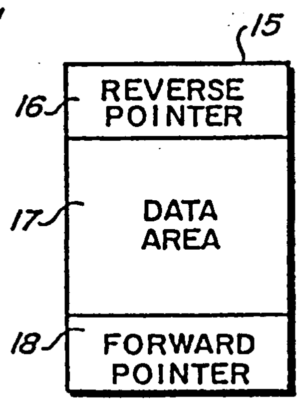

FIG. 2

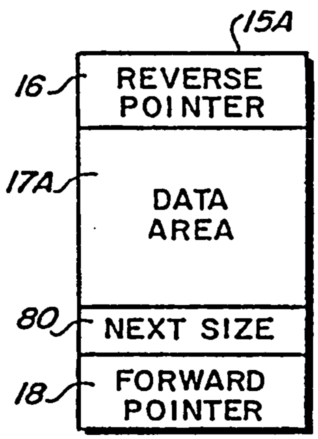

FIG. 6

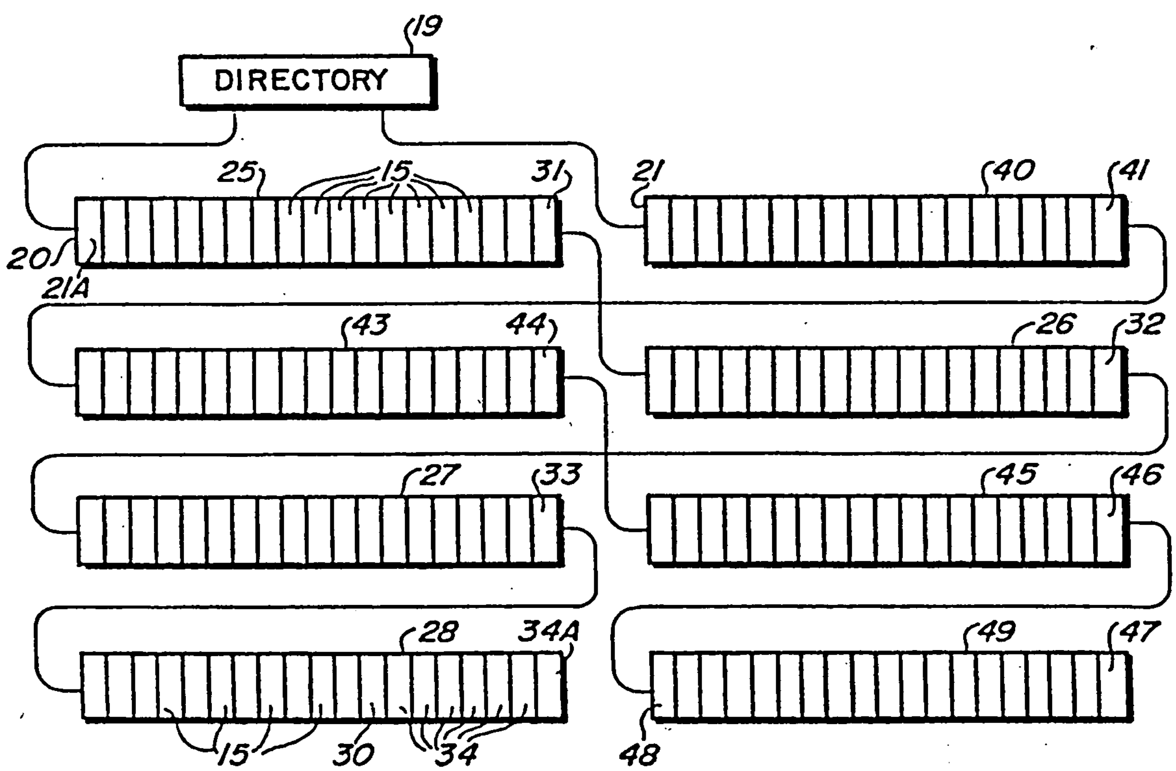

FIG. 3

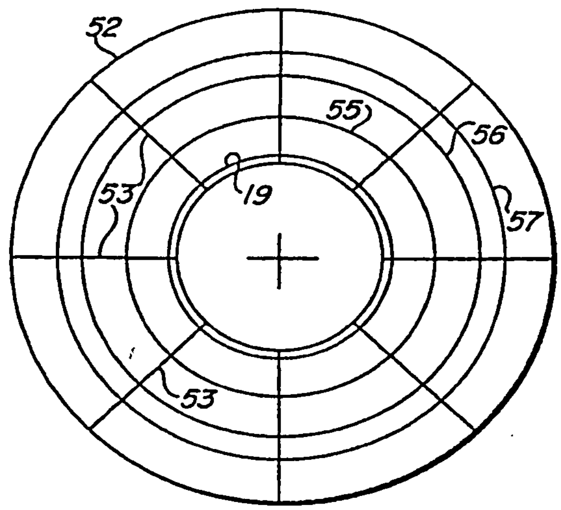
52
55
56
53
19
57
53

FIG. 4

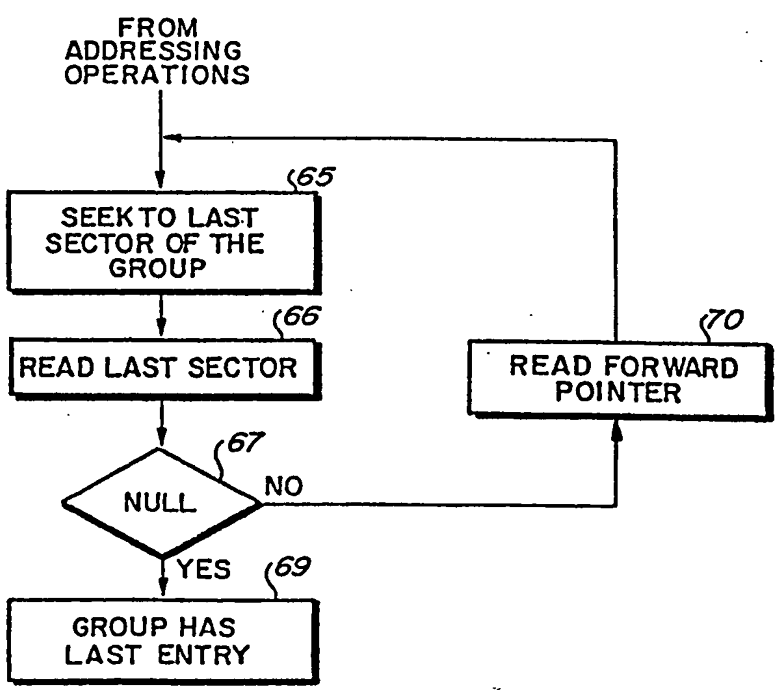
FROM ADDRESSING OPERATIONS
65
SEEK TO LAST SECTOR OF THE GROUP
66
READ LAST SECTOR
70
READ FORWARD POINTER
67
NULL
NO
YES
69
GROUP HAS LAST ENTRY

FIG. 5